


NASA/CR-2002-211759
ICASE Report No. 2002-26



A Logical Process Calculus

Rance Cleaveland

State University of New York at Stony Brook, Stony Brook, New York

Gerald Lüttgen

The University of Sheffield, Sheffield, United Kingdom



DISTRIBUTION STATEMENT A
Approved for Public Release
Distribution Unlimited

August 2002

20021017 108

The NASA STI Program Office . . . in Profile

Since its founding, NASA has been dedicated to the advancement of aeronautics and space science. The NASA Scientific and Technical Information (STI) Program Office plays a key part in helping NASA maintain this important role.

The NASA STI Program Office is operated by Langley Research Center, the lead center for NASA's scientific and technical information. The NASA STI Program Office provides access to the NASA STI Database, the largest collection of aeronautical and space science STI in the world. The Program Office is also NASA's institutional mechanism for disseminating the results of its research and development activities. These results are published by NASA in the NASA STI Report Series, which includes the following report types:


- **TECHNICAL PUBLICATION.** Reports of completed research or a major significant phase of research that present the results of NASA programs and include extensive data or theoretical analysis. Includes compilations of significant scientific and technical data and information deemed to be of continuing reference value. NASA's counterpart of peer-reviewed formal professional papers, but having less stringent limitations on manuscript length and extent of graphic presentations.
- **TECHNICAL MEMORANDUM.** Scientific and technical findings that are preliminary or of specialized interest, e.g., quick release reports, working papers, and bibliographies that contain minimal annotation. Does not contain extensive analysis.
- **CONTRACTOR REPORT.** Scientific and technical findings by NASA-sponsored contractors and grantees.
- **CONFERENCE PUBLICATIONS.** Collected papers from scientific and technical conferences, symposia, seminars, or other meetings sponsored or cosponsored by NASA.
- **SPECIAL PUBLICATION.** Scientific, technical, or historical information from NASA programs, projects, and missions, often concerned with subjects having substantial public interest.
- **TECHNICAL TRANSLATION.** English-language translations of foreign scientific and technical material pertinent to NASA's mission.

Specialized services that complement the STI Program Office's diverse offerings include creating custom thesauri, building customized data bases, organizing and publishing research results . . . even providing videos.

For more information about the NASA STI Program Office, see the following:

- Access the NASA STI Program Home Page at <http://www.sti.nasa.gov>
- Email your question via the Internet to help@sti.nasa.gov
- Fax your question to the NASA STI Help Desk at (301) 621-0134
- Telephone the NASA STI Help Desk at (301) 621-0390
- Write to:
NASA STI Help Desk
NASA Center for AeroSpace Information
7121 Standard Drive
Hanover, MD 21076-1320

NASA/CR-2002-211759
ICASE Report No. 2002-26



A Logical Process Calculus

Rance Cleaveland

State University of New York at Stony Brook, Stony Brook, New York

Gerald Lüttgen

The University of Sheffield, Sheffield, United Kingdom

ICASE

NASA Langley Research Center

Hampton, Virginia

Operated by Universities Space Research Association



Prepared for Langley Research Center
under Contract NAS1-97046

August 2002

Available from the following:

NASA Center for AeroSpace Information (CASI)
7121 Standard Drive
Hanover, MD 21076-1320
(301) 621-0390

National Technical Information Service (NTIS)
5285 Port Royal Road
Springfield, VA 22161-2171
(703) 487-4650

A LOGICAL PROCESS CALCULUS*

RANCE CLEAVELAND[†] AND GERALD LÜTTGEN[‡]

Abstract. This paper presents the Logical Process Calculus (LPC), a formalism that supports *heterogeneous* system specifications containing both operational and declarative subspecifications. Syntactically, LPC extends Milner's Calculus of Communicating Systems with operators from the alternation-free linear-time μ -calculus ($\text{LT}\mu$). Semantically, LPC is equipped with a behavioral preorder that generalizes Hennessy's and DeNicola's must-testing preorder as well as $\text{LT}\mu$'s satisfaction relation, while being compositional for all LPC operators. From a technical point of view, the new calculus is distinguished by the inclusion of (i) both minimal and maximal fixed-point operators and (ii) an unimplementability predicate on process terms, which tags inconsistent specifications. The utility of LPC is demonstrated by means of an example highlighting the benefits of heterogeneous system specification.

Key words. heterogeneous specification, must-testing, process algebra, temporal logic, testing theory

Subject classification. Computer Science

1. Introduction. Over the past two decades, a wealth of approaches to formally specifying and reasoning about reactive systems have been introduced. Most of these may be classified according to whether they are based on *process algebra* [3] or *temporal logic* [27]. The process-algebraic paradigm is founded on notions of *refinement*, where one typically formulates a system specification and its implementation in the same notation and then proves that the latter refines the former. The underlying semantics is usually given operationally, and refinement relations are formalized as preorders. In contrast, the temporal-logic paradigm is based on the use of *temporal logics* [27] to formulate specifications, with implementations being given in an operational notation. One then verifies a system by establishing that it is a *model* of its specification, in the formal logical sense. The strength of the former paradigm is its support for *compositional reasoning*, i.e., one may refine system components independently of others. The benefit of the latter paradigm originates in its support for abstract specifications, where irrelevant operational details may be ignored. Both approaches may be given automated support in the form of *model checking* when the considered systems are finite-state.

The objective of this paper is to develop a *compositional theory for heterogeneous specifications* that uniformly integrates both refinement-based and temporal-logic specification styles, thereby allowing both approaches to be taken advantage of when designing systems. Accordingly, we present a novel *Logical Process Calculus* (LPC) that combines the algebraic operators of Milner's *Calculus of Communicating Systems* (CCS) [25] with the logical operators of the *Alternation-Free Linear-Time μ -Calculus* ($\text{LT}\mu$) [32]. More precisely, we show that logical disjunction in $\text{LT}\mu$ may be understood as internal choice, complementing

*This work was supported by the National Aeronautics and Space Administration under NASA Contract No. NAS1-97046 while the second author was in residence at ICASE, NASA Langley Research Center, Hampton, Virginia 23681-2199, USA. Research support was also provided by the National Science Foundation under NSF grant CCR-9988489.

[†]Department of Computer Science, State University of New York at Stony Brook, Stony Brook, New York 11794-4400, USA, e-mail: rance@cs.sunysb.edu.

[‡]Department of Computer Science, The University of Sheffield, 211 Portobello Street, Sheffield S1 4DP, U.K., e-mail: g.luetngen@dcs.shef.ac.uk.

the external choice operator in CCS, and logical conjunction in $\text{LT}\mu$ as synchronous parallel composition, complementing asynchronous parallel composition in CCS. Moreover, $\text{LT}\mu$ is equipped with two recursion operators, a least fixed-point operator and a greatest fixed-point operator, which allow for the finite but unbounded, and the infinite, unwinding of recursion, respectively. The behavior described by the greatest fixed-point operator in $\text{LT}\mu$ thus corresponds to recursion in CCS. In the light of this discussion, LPC extends CCS by operators for *disjunction*, *conjunction*, and *minimal fixed-points*, as well as the basic processes *true* and *false*, and thereby allows for the encoding of both $\text{LT}\mu$ formulas and CCS processes in LPC (cf. Sec. 2).

The semantics of LPC is based on the testing approach of DeNicola and Hennessy [11]. The hallmarks of this theory are the use of transitions to model both processes and tests and the differentiation of processes on the basis of their responses to tests. Accordingly, we equip LPC terms with a transition relation defining the single-step transitions that specifications may engage in. We also introduce a novel *unimplementability predicate* on terms whose role is to identify inconsistent specifications, such as *false*, that cannot be implemented. Both the transition relation and the unimplementability predicate are defined via structural operational rules, i.e., in a syntax-driven fashion. We then carry over the definitions of must-testing in [11] to our setting and show that the resulting behavioral preorder (i) conservatively extends the traditional must-preorder between CCS specifications; (ii) is compositional for all operators in LPC; and (iii) naturally encodes the standard satisfaction relation between CCS processes and $\text{LT}\mu$ formulas (cf. Sec. 3). Thus, our framework may be seen to unify refinement-based and logic-based approaches to system specification, while facilitating component-based reasoning. Technically, this expressiveness follows from the mathematically coherent inclusion of process and logical operators in LPC that is enabled by our treatment of unimplementability (cf. Sec. 4). Practically, the theory allows system modelers to freely intermix operational and declarative subspecifications using both system operators (e.g. parallel composition) and logical constructors (e.g. conjunction). This gives engineers powerful tools to model system components at different levels of abstraction and to impose declarative constraints on the execution behavior of components (cf. Sec. 5).

2. A Logical Process Calculus. This section formally introduces our logical process calculus, LPC. We first present its syntax and then define its semantics via operational rules and a novel unimplementability predicate. Finally, the calculus is equipped with a refinement preorder on processes, which is an adaptation of DeNicola and Hennessy’s must-testing preorder [11].

2.1. Syntax of LPC. The syntax of LPC extends Milner’s CCS [25] with *disjunction*, *conjunction*, and *least fixed-point* operators. It also includes a process constant for the universal process *true*, while *false* will be a derived process term in our calculus. Formally, let Λ be a countable set of *actions*, or ports, not including the distinguished unobservable, *internal* action τ . With every $a \in \Lambda$ we associate a *complementary action* \bar{a} . We define $\bar{\Lambda} := \{\bar{a} \mid a \in \Lambda\}$ and take \mathcal{A} to denote the set $\Lambda \cup \bar{\Lambda}$. Complementation is lifted to \mathcal{A} by defining $\bar{\bar{a}} := a$. As in CCS, an action a communicates with its complement \bar{a} to produce the internal action τ . We let a, b, \dots range over \mathcal{A} and α, β, \dots over $\mathcal{A}_\tau := \mathcal{A} \cup \{\tau\}$. The syntax of LPC is then defined as follows:

$$\begin{aligned} P ::= & \mathbf{0} \mid \mathbf{tt} \mid x \mid w \mid \alpha.P \mid P + P \mid P \vee P \mid P|P \mid P \wedge P \mid \\ & P \setminus L \mid P[f] \mid \mu x.P \mid \mu_k x.P \mid \nu x.P \end{aligned}$$

where $k \in \mathbb{N}$, x is a *variable* taken from some nonempty set \mathcal{V} of variables, w is an infinite word over \mathcal{A} whose inclusion will be discussed in the next section, set $L \subseteq \mathcal{A}$ is a *restriction set*, and $f : \mathcal{A}_\tau \rightarrow \mathcal{A}_\tau$ is a *finite relabeling*. A finite relabeling satisfies the properties $f(\tau) = \tau$, $f(\bar{a}) = \overline{f(a)}$, and $|\{\alpha \mid f(\alpha) \neq \alpha\}| < \infty$. We

define $\bar{L} := \{\bar{a} \mid a \in L\}$ and use the standard definitions for *free* and *bound* variables, *open* and *closed* terms, *guardedness*, and *contexts*. We require for fixed-point terms $\mu x.P$, $\mu_k x.P$, and $\nu x.P$ that x is guarded in P . Intuitively, $\mu x.P$ stands for finite *unbounded* unwindings of P , while $\mu_k x.P$ encodes finite unwindings of P *bounded* by k . A term is called *alternation-free* if every variable bound by a least (greatest) fixed-point $\mu x.P$ ($\nu x.P$) does not occur free in a subterm $\nu y.Q$ ($\mu y.Q$) of P . We refer to closed, guarded, and alternation-free¹ terms as *processes*, with the set of all processes written as \mathcal{P} . Finally, we denote syntactic equality by \equiv .

While it is obvious that LPC subsumes all CCS processes, it is not immediately clear that it also encodes all Alternation-Free Linear-Time μ -Calculus (LT μ) formulas [5]². The syntax of LT μ formulas is given by the following BNF:

$$\Phi ::= 0 \mid \text{tt} \mid \text{ff} \mid x \mid \langle a \rangle \Phi \mid \Phi \vee \Phi \mid \Phi \wedge \Phi \mid \mu x. \Phi \mid \nu x. \Phi$$

In our setting, LT μ formulas will be interpreted over infinite action sequences and also finite ones leading to deadlock. This is why the ‘deadlock formula’ 0 is included in LT μ . In LPC, ff corresponds to the term $\mu x. \tau. x$ and the *next* operator ‘ $\langle a \rangle$ ’, for $a \in \mathcal{A}$, to the prefix operator ‘ $a.$ ’.

2.2. Semantics of LPC. The *operational semantics* of an LPC process P is given as a labeled transition system $\langle \mathcal{P}, \mathcal{A}_\tau, \longrightarrow, \#, P \rangle$, where \mathcal{P} is the set of states, \mathcal{A}_τ the alphabet, $\longrightarrow \subseteq \mathcal{P} \times \mathcal{A}_\tau \times \mathcal{P}$ the transition relation, $\# \subseteq \mathcal{P}$ our *unimplementability predicate* that is discussed below, and P the start state.

The transition relation is defined by the *structural operational rules* displayed in Table 2.1. For convenience, we write $P \xrightarrow{\alpha} P'$ instead of $\langle P, \alpha, P' \rangle \in \longrightarrow$. Note that, for the CCS operators, the semantics is exactly as in [25]. As for the other constructs, tt can nondeterministically engage in any action transition, or decide to deadlock (cf. Rules (True1) and (True2)). Process $\alpha.P$ may engage in action α and then behave like P (cf. Rule (Act1)), and similarly the process described by the infinite word aw may engage in its initial action a and then behave like w (cf. Rule (Act2)). The reason for including process w is to enable the modeling of arbitrary system environments within our calculus, including those exhibiting irregular behavior. The summation operator $+$ denotes *nondeterministic external choice* such that $P + Q$ may behave like P or Q , depending on which communication initially offered by P and Q is accepted by the environment (cf. Rules (Sum1) and (Sum2)). Analogously, \vee encodes *disjunction* or *nondeterministic internal choice*, i.e., process $P \vee Q$ determines internally, without consulting its environment, whether to execute P or Q (cf. Rules (Dis1) and (Dis2)). Process $P|Q$ stands for the *asynchronous parallel composition* of processes P and Q according to an interleaving semantics with synchronized communication on complementary actions, resulting in the internal action τ (cf. Rules (Par1)–(Par3)). Similarly, $P \wedge Q$ encodes the *conjunction* or *synchronous parallel composition* of P and Q , with synchronization on all visible actions and interleaving on τ (cf. Rules (Con1)–(Con3)). The *restriction operator* $\backslash L$ prohibits the execution of actions in $L \cup \bar{L}$ and, thus, permits the scoping of actions. Process $P[f]$ behaves exactly as P where actions are renamed according to the *relabeling* f . The remaining rules define the semantics of our least and greatest fixed-point operators. The *minimal fixed-point* process $\mu x.P$ first guesses some number $k \in \mathbb{N}$ that determines how often P might be unwound, as encoded by the process $\mu_k x.P$ (cf. Rules (Mu1) and (Mu2))³. Here, $P[Q/x]$ stands for the process P with all of its free occurrences of variable x substituted by Q . This account of μ

¹The restriction to alternation-free processes is made for continuity reasons that are elaborated on later.

²LT μ is more expressive than linear-time temporal logic, so the limitation to alternation-free formulas does not impose undue expressiveness restrictions.

³The presence of unbounded internal choice in Rules (True1) and (Mu1) presents problems for more denotational process theories; in LPC it proves not to be problematic because of our exclusively operational orientation.

TABLE 2.1
Operational semantics

True1 $\frac{}{\text{tt} \xrightarrow{\tau} a.\text{tt}} a \in \mathcal{A}$	True2 $\frac{}{\text{tt} \xrightarrow{\tau} \mathbf{0}}$
Act1 $\frac{}{\alpha.P \xrightarrow{\alpha} P}$	Act2 $\frac{}{aw \xrightarrow{a} w}$
Sum1 $\frac{P \xrightarrow{\alpha} P'}{P + Q \xrightarrow{\alpha} P'}$	Sum2 $\frac{Q \xrightarrow{\alpha} Q'}{P + Q \xrightarrow{\alpha} Q'}$
Dis1 $\frac{}{P \vee Q \xrightarrow{\tau} P}$	Dis2 $\frac{}{P \vee Q \xrightarrow{\tau} Q}$
Par1 $\frac{P \xrightarrow{\alpha} P'}{P Q \xrightarrow{\alpha} P' Q}$	Par2 $\frac{Q \xrightarrow{\alpha} Q'}{P Q \xrightarrow{\alpha} P Q'}$
Con1 $\frac{P \xrightarrow{\tau} P'}{P \wedge Q \xrightarrow{\tau} P' \wedge Q}$	Con2 $\frac{Q \xrightarrow{\tau} Q'}{P \wedge Q \xrightarrow{\tau} P \wedge Q'}$
Par3 $\frac{P \xrightarrow{a} P' \quad Q \xrightarrow{\bar{a}} Q'}{P Q \xrightarrow{\tau} P' Q'}$	Con3 $\frac{P \xrightarrow{a} P' \quad Q \xrightarrow{a} Q'}{P \wedge Q \xrightarrow{a} P' \wedge Q'}$
Res $\frac{P \xrightarrow{\alpha} P'}{P \setminus L \xrightarrow{\alpha} P' \setminus L} \alpha \notin L \cup \bar{L}$	Rel $\frac{P \xrightarrow{\alpha} P'}{P[f] \xrightarrow{f(\alpha)} P'[f]}$
Mu1 $\frac{}{\mu x.P \xrightarrow{\tau} \mu_k x.P} k \in \mathbb{N}$	Mu2 $\frac{P[\mu_{k-1} x.P/x] \xrightarrow{\alpha} P'}{\mu_k x.P \xrightarrow{\alpha} P'} k > 0$
Nu $\frac{P[\nu x.P/x] \xrightarrow{\alpha} P'}{\nu x.P \xrightarrow{\alpha} P'}$	

may be seen as embodying a form of *continuity*: μ is interpreted in terms of its finite unwindings. Because of continuity problems associated with alternating least and greatest fixed points, in this paper we only consider alternation-free process expressions. The *maximal fixed-point* process $\nu x.P$ may unwind its loop indefinitely, as is the case for recursion in CCS (cf. Rule (Nu)). Note that the purely *divergent process* Ω , employed in some process algebras [16] for describing infinite internal computation, can be derived in LPC as $\nu x.\tau.x$.

Temporal logics, including $\text{LT}\mu$, are capable of specifying *inconsistencies* or contradictions, i.e., behaviors equivalent to *false*. From an operational point of view, a process describing an inconsistency is not implementable, and thus runs of processes passing through unimplementable states should be ignored. Due to nondeterministic choice, a process that can engage in such runs is not necessarily unimplementable itself. It is only unimplementable if all of its runs must pass through an unimplementable state. This intuition is reflected in the definition of our *unimplementability predicate*, given in Table 2.2, where we write $P \#$ for

TABLE 2.2
Unimplementability predicate #

-
1. $\mu_0 x.P \#$
 2. $P \longrightarrow$ and $P \wedge Q \not\rightarrow$ implies $P \wedge Q \#$
 3. $Q \longrightarrow$ and $P \wedge Q \not\rightarrow$ implies $P \wedge Q \#$
 4. $P \#$ implies
 - $\alpha.P \#$ • $P[f] \#$ • $P \setminus L \#$
 - $P \wedge Q \#$ • $Q \wedge P \#$
 - $P|Q \#$ • $Q|P \#$
 - $\nu x.P \#$ • $\mu x.P \#$ • $\mu_k x.P \#$
 5. $P \#$ and $Q \#$ implies
 - $P + Q \#$ • $P \vee Q \#$
 6. $P[\mu_{k-1} x.P/x] \#$ implies $\mu_k x.P \#$, for $k > 0$
 7. $(\forall k. \mu_k x.P \#)$ implies $\mu x.P \#$
-

$P \in \#$ and where $P \longrightarrow$ stands for $\exists P' \in \mathcal{P} \exists \alpha \in \mathcal{A}_\tau. P \xrightarrow{\alpha} P'$. In particular, a contradiction is present within a conjunction $P \wedge Q$, if the conjunction process cannot engage in any transition, although one of its argument processes can (cf. Rules (2) and (3)). As an example, consider process $a.0 \wedge b.0$, for $a \neq b$. Further, Rule (1) states that the unimplementability of P propagates backwards through prefixing. Note that the operational semantics for LPC distinguishes between inconsistent processes that are unimplementable and deadlocked processes that are implementable. For example, both processes $(a.0|b.0) \setminus \{a, b\}$ and $a.0 \wedge b.0$ cannot engage in any transitions. However, $(a.0 \wedge b.0) \#$ while $\neg(((a.0|b.0) \setminus \{a, b\}) \#)$, as desired. All other rules are straightforward, except for least fixed-point processes, such as the process $\mu_0 x.P$ that cannot unwind its body P further and is thus considered to be unimplementable (cf. Rule (1)). Together with Rules (6) and (7), this implies that the process $\mu x.\tau.x$, which can engage in finite but unbounded numbers of τ 's, is actually unimplementable. Indeed, we will identify this process with *false* and abbreviate it by *ff*. Finally, it is easy to prove via induction on the structure of process terms that $P \xrightarrow{\alpha} P'$ and $P \#$ implies $P' \#$, for any $P, P' \in \mathcal{P}$ and $\alpha \in \mathcal{A}_\tau$.

The semantics for LPC does not only extend the standard CCS semantics but is also compatible with the semantics of $\text{LT}\mu$ formulas; see Thm. 3.5. This theorem, however, is not straightforward, and its proof requires us to build a rich semantic theory for LPC. Before doing so we first introduce some notation. A *potential path* π of process P is a sequence of transitions $(P_i \xrightarrow{\alpha_i} P_{i+1})_{0 \leq i < k}$, for some $k \in \mathbb{N} \cup \{\omega\}$, such that $P_0 \equiv P$. If $\neg(P_i \#)$, for all $0 \leq i < k$, then π is called an *implementable path*, or simply *path*. We use $|\pi|$ to refer to k , the length of π . If $|\pi| = \omega$, we say that π is *infinite*; otherwise, π is *finite*. Moreover, π is called *maximal* if $|\pi| < \omega$ and $P_{|\pi|} \not\rightarrow$. The *trace* $\text{trace}(\pi)$ of π is defined as the word $w := (\alpha_i)_{I_\pi} \in \mathcal{A}^\omega := \mathcal{A}^* \cup \mathcal{A}^\omega$, where $I_\pi := \{0 \leq i < |\pi| \mid \alpha_i \neq \tau\}$. In the case of $I_\pi = \emptyset$, we let ϵ stand for $w = ()$. Moreover, if π is finite, we also write $P \xRightarrow{w} P_{|\pi|}$ for π . We denote the sets of all finite, maximal, and infinite paths of P by $\Pi_{\text{fin}}(P)$, $\Pi_{\text{max}}(P)$, and $\Pi_\omega(P)$, respectively. We may also introduce according languages for P :

$$\begin{aligned}
 \mathcal{L}_{\text{fin}}(P) &:= \{\text{trace}(\pi) \mid \pi \in \Pi_{\text{fin}}(P)\} &\subseteq \mathcal{A}^* &\quad \text{finite-trace language of } P \\
 \mathcal{L}_{\text{max}}(P) &:= \{\text{trace}(\pi) \mid \pi \in \Pi_{\text{max}}(P)\} &\subseteq \mathcal{A}^* &\quad \text{maximal-trace language of } P \\
 \mathcal{L}_\omega(P) &:= \{\text{trace}(\pi) \mid \pi \in \Pi_\omega(P)\} &\subseteq \mathcal{A}^\omega &\quad \text{infinite-trace language of } P
 \end{aligned}$$

The semantic theory to be developed for LPC relies on the notion of *divergence*, i.e., a system's ability to engage in an infinite internal computation. In this paper, we employ the traditional notion of divergence

as used by DeNicola and Hennessy [11]; more sophisticated definitions may be found elsewhere in the literature [6, 26, 28]. A process P is *divergent*, in signs $P \uparrow$, if $\epsilon \in \mathcal{L}_\omega(P)$. For example, the process $\Omega := \nu x. \tau. x$, is divergent. A process P is called *w-divergent* for some $w \in \mathcal{A}^\infty$, in signs $P \uparrow w$, if $\exists P' \in \mathcal{P} \exists v <_{\text{fin}} w. P \xRightarrow{v} P'$ and $P' \uparrow$. Here, $<_{\text{fin}}$ stands for the prefix ordering on words. We further write $\mathcal{L}_{\text{div}}(P)$ for the *divergent-trace language* of P , i.e., $\mathcal{L}_{\text{div}}(P) := \{w \in \mathcal{A}^\infty \mid P \uparrow w\}$. Finally, P is called *convergent* or *w-convergent*, in symbols $P \downarrow$ and $P \downarrow w$, if $\neg(P \uparrow)$ and $\neg(P \uparrow w)$, respectively.

2.3. Refinement in LPC. We now turn our attention to a behavioral theory of LPC, which defines a behavioral preorder \sqsubseteq on processes such that $P \sqsubseteq Q$, i.e., Q refines P , if Q is “more defined” than P . The preorder is an adaptation of DeNicola and Hennessy’s *must-preorder* [11], which was developed within an elegant *testing theory* and distinguishes processes on the basis of the *tests* they are necessarily able to pass. In this context, tests are processes equipped with a special action \surd , which are employed to witness the interactions a process may have with its environment. In order to determine whether a process passes a test, one has to examine the maximal and infinite *computations* that result when the test runs in lock-step with the process under consideration.

Formally, a *test* is a process that might use the distinguished success action $\surd \notin \mathcal{A}_\tau$. The set of all tests is denoted by \mathcal{T} . A *maximal (infinite) computation* π of process P and test T is a maximal (infinite) path π of $(P|T) \setminus \mathcal{A}$, i.e., $\pi = ((P_i|T_i) \setminus \mathcal{A} \xrightarrow{\tau} (P_{i+1}|T_{i+1}) \setminus \mathcal{A})_{0 \leq i < |\pi|}$. Recall that paths only go along implementable states. Computation π is *successful* if $T_k \xrightarrow{\surd}$ for some $0 \leq k < |\pi|$; otherwise, it is *unsuccessful*. Finally, process P is said to *must-satisfy* test T , in symbols $P \text{ must } T$, if every maximal and infinite computation of P and T is successful. Our variant of the must-preorder can now be defined as follows.

DEFINITION 2.1 (Must-preorder). *For $P, Q \in \mathcal{P}$ we let $P \sqsubseteq Q$ if, for all $T \in \mathcal{T}$, $P \text{ must } T$ implies $Q \text{ must } T$.*

It is easy to see that \sqsubseteq is a preorder, i.e., that it is reflexive and transitive. Note that this preorder can be extended to open terms by the usual means of closed substitution [25]. Moreover, \sqsubseteq satisfies the following basic algebraic laws, where \approx stands for the kernel $\sqsubseteq \cap (\sqsubseteq)^{-1}$ of \sqsubseteq .

PROPOSITION 2.2. *Let $P, Q, R \in \mathcal{P}$. Then, the following holds:*

$$\begin{array}{llll} P|Q \approx Q|P & (P|Q)|R \approx P|(Q|R) & P|0 \approx P & P|\Omega \approx \Omega \\ P \wedge Q \approx Q \wedge P & (P \wedge Q) \wedge R \approx P \wedge (Q \wedge R) & P \wedge tt \approx P & P \wedge ff \approx ff \\ P + Q \approx Q + P & (P + Q) + R \approx P + (Q + R) & P + 0 \approx P & P + \Omega \approx \Omega \\ P \vee Q \approx Q \vee P & (P \vee Q) \vee R \approx P \vee (Q \vee R) & P \vee tt \approx tt & P \vee ff \approx P \end{array}$$

Further, $P \wedge P \approx P$, $P \vee P \approx P$, and $P \vee Q \sqsubseteq P$.

It is also easy to see that the divergent process Ω does not must-satisfy any tests, except the trivial ones, such as $\surd.0$. Hence, it is the smallest process with respect to \sqsubseteq . Conversely, process ff must-satisfies every test, since it does not possess any computation due to $ff\#$. Consequently, ff is the largest process with respect to \sqsubseteq . Also tt is a distinguished process in our setting; it is the smallest *convergent* process with respect to \sqsubseteq . Thus, we have $\Omega \sqsubseteq tt \sqsubseteq ff$ ⁴.

⁴This ordering is the reverse of the more usual Boolean ordering, which holds that ff is lower than tt . Our ordering is due to the fact that must refinement implies *reverse* language containment.

3. Properties of the Must-Preorder. In this section we investigate the utility of our calculus for the heterogeneous specification of reactive systems. We show that our must-preorder is a conservative extension of the one of DeNicola and Hennessy, provide its characterization in terms of traces and initial action sets, investigate its close relation to $\text{LT}\mu$ satisfaction, and finally establish its compositionality properties.

3.1. Extension of DeNicola and Hennessy's Must-Preorder. It is easy to see that our must-preorder \preceq is a conservative extension of the original must-preorder \preceq_{DH} of DeNicola and Hennessy, defined on CCS processes [11]. The reason is that their and our definitions of the testing framework coincide on CCS processes. Hence, we may formally obtain the following conservativity theorem.

THEOREM 3.1. *Let P, Q be CCS processes. Then, $P \preceq Q$ if and only if $P \preceq_{\text{DH}} Q$.*

3.2. Characterization. We now present a characterization of our must-preorder which will be used for obtaining some of our main results. The characterization closely follows the lines of a similar characterization of DeNicola and Hennessy's must-preorder [11]. It uses the notation $\mathcal{I}(P)$ for the set $\{a \in \mathcal{A} \mid P \xrightarrow{\tau}^* \xrightarrow{a}\}$ of visible initial actions of P .

THEOREM 3.2. *Let P, Q be processes. Then $P \preceq Q$ if and only if for all $w \in \mathcal{A}^\infty$ such that $P \Downarrow w$:*

1. $Q \Downarrow w$
2. $|w| < \omega$: $\forall Q'. Q \xRightarrow{w} Q'$ implies $\exists P'. P \xRightarrow{w} P'$ and $\mathcal{I}(P') \subseteq \mathcal{I}(Q')$
3. $|w| = \omega$: $w \in \mathcal{L}_\omega(Q)$ implies $w \in \mathcal{L}_\omega(P)$

Observe that this characterization is also sensitive to infinite traces and not only finite ones (cf. Cond. (2)). This is superficially similar to the *improved failures model* of [7]; the difference is that infinite traces in [7] convey divergence information, while they convey convergence information in the above characterization.

The proof of the above theorem relies on the following four distinguished tests, where $k \in \mathbb{N}$, $w = (a_i)_{0 \leq i < k} \in \mathcal{A}^*$, $v \in \mathcal{A}^\omega$, and $a \in \mathcal{A}$.

1. $T_w^\Downarrow := a_0.a_1 \cdots a_{k-1}.0 \mid \tau.\sqrt{.0}$
2. $T_w^{\text{fin}} := a_0.(a_1 \cdots (a_{k-1}.0 + \tau.\sqrt{.0}) \cdots) + \tau.\sqrt{.0} + \tau.\sqrt{.0}$
3. $T_{w,a}^{\text{max}} := a_0.(a_1 \cdots (a_{k-1}.a.\sqrt{.0} + \tau.\sqrt{.0}) \cdots) + \tau.\sqrt{.0} + \tau.\sqrt{.0}$
4. $T_v^\omega := v \mid \tau.\sqrt{.0}$

The intuitions behind defining these tests are as follows.

LEMMA 3.3. *Let P be an arbitrary LPC process and*

1. *Let $w \in \mathcal{A}^*$. Then, $P \Downarrow w$ iff $P \text{ must } T_w^\Downarrow$.*
2. *Let $w \in \mathcal{A}^*$ such that $P \Downarrow w$. Then, $w \notin \mathcal{L}_{\text{fin}}(P)$ iff $P \text{ must } T_w^{\text{fin}}$.*
3. *Let $w \in \mathcal{A}^*$ such that $P \Downarrow w$. Then, $w \notin \mathcal{L}_{\text{max}}(P)$ iff $\exists a \in \mathcal{A}. P \text{ must } T_{w,a}^{\text{max}}$.*
4. *Let $v \in \mathcal{A}^\omega$ such that $P \Downarrow v$. Then, $v \notin \mathcal{L}_\omega(P)$ iff $P \text{ must } T_v^\omega$.*

The proof of this lemma is not too difficult but tedious; it follows our definition of must-passing tests and is similar to the corresponding proof in [9]. Note that the first property can also be carried over to infinite words, due to our 'approximative' definition of divergence.

3.3. Extension of $\text{LT}\mu$ Satisfaction. To prove that our must-preorder is also an extension of $\text{LT}\mu$ satisfaction we first recall the standard semantics of $\text{LT}\mu$. An $\text{LT}\mu$ formula is interpreted as the set of those finite and infinite sequences over \mathcal{A} that validate the formula. Formally, the semantics $[\Phi]^\varepsilon$ of a possibly

open $\text{LT}\mu$ term Φ is defined relative to an environment \mathcal{E} mapping variables to subsets of \mathcal{A}^∞ . Note that our variant of the linear-time μ -calculus [5] can be used to reason about deadlock traces as well, due to our inclusion of the atomic proposition $\mathbf{0}$; this is why we also consider finite traces, in addition to infinite ones.

$$\begin{aligned} \llbracket \text{tt} \rrbracket^\mathcal{E} &:= \mathcal{A}^\infty & \llbracket \text{ff} \rrbracket^\mathcal{E} &:= \emptyset & \llbracket x \rrbracket^\mathcal{E} &:= \mathcal{E}(x) \\ \llbracket \langle a \rangle \Phi \rrbracket^\mathcal{E} &:= \{aw \mid w \in \llbracket \Phi \rrbracket^\mathcal{E}\} & \llbracket \mathbf{0} \rrbracket^\mathcal{E} &:= \{\epsilon\} \\ \llbracket \mu x. \Phi \rrbracket^\mathcal{E} &:= \bigcap \{T \subseteq \mathcal{A}^\infty \mid \llbracket \Phi \rrbracket^{\mathcal{E}[x \mapsto T]} \subseteq T\} & \llbracket \Phi_1 \wedge \Phi_2 \rrbracket^\mathcal{E} &:= \llbracket \Phi_1 \rrbracket^\mathcal{E} \cap \llbracket \Phi_2 \rrbracket^\mathcal{E} \\ \llbracket \nu x. \Phi \rrbracket^\mathcal{E} &:= \bigcup \{T \subseteq \mathcal{A}^\infty \mid T \subseteq \llbracket \Phi \rrbracket^{\mathcal{E}[x \mapsto T]}\} & \llbracket \Phi_1 \vee \Phi_2 \rrbracket^\mathcal{E} &:= \llbracket \Phi_1 \rrbracket^\mathcal{E} \cup \llbracket \Phi_2 \rrbracket^\mathcal{E} \end{aligned}$$

In case Φ is a formula, i.e., Φ is a closed $\text{LT}\mu$ term, it is easy to see that the environment \mathcal{E} is irrelevant. We say that a CCS process P satisfies Φ , in signs $P \models \Phi$, if all traces of P are included in the traces of $\llbracket \Phi \rrbracket$. Formally, $P \models \Phi$ if (i) $\mathcal{L}_{\text{div}}(P) \subseteq \mathcal{L}_{\text{div}}(\Phi)$, (ii) $\mathcal{L}_{\text{max}}(P) \subseteq \llbracket \Phi \rrbracket$, and (iii) $\mathcal{L}_\omega(P) \subseteq \llbracket \Phi \rrbracket$.

Further, $\text{LT}\mu$ formulas, when considered as a sublanguage of LPC, possess two important properties. First, all formulas Φ are convergent, i.e., $\mathcal{L}_{\text{div}}(\Phi) = \emptyset$. This is because the internal prefix operator ‘ τ .’ is not available in $\text{LT}\mu$. In addition, the atomic propositions tt , ff , and $\mathbf{0}$ do not give rise to divergence. As a consequence, Cond. (i) in the definition of $P \models \Phi$ above can be simplified to $\mathcal{L}_{\text{div}}(P) = \emptyset$. In particular, formula tt is satisfied by convergent processes only, whence $P \models \text{tt}$ if and only if $\mathcal{L}_{\text{div}}(P) = \emptyset$. Second, every $\text{LT}\mu$ formula Φ is purely nondeterministic in the sense that all choices are internal:

$$\forall \Phi', \Phi'' \forall \alpha, \beta. \Phi \xrightarrow{\alpha} \Phi', \Phi \xrightarrow{\beta} \Phi'', \Phi' \neq \Phi'' \text{ implies } \alpha \equiv \beta \equiv \tau.$$

This is due to the fact that disjunction is modeled as internal choice in LPC.

PROPOSITION 3.4. *Let Φ be an $\text{LT}\mu$ formula and P a CCS process. Then, $\Phi \preceq P$ if and only if (i) $\mathcal{L}_{\text{div}}(P) = \emptyset$, (ii) $\mathcal{L}_{\text{max}}(P) \subseteq \mathcal{L}_{\text{max}}(\Phi)$, and (iii) $\mathcal{L}_\omega(P) \subseteq \mathcal{L}_\omega(\Phi)$.*

The proof of this proposition relies on our characterization theorem for \preceq (cf. Thm. 3.2) and uses the two properties of formulas mentioned above. The proposition is the key for establishing the next theorem.

THEOREM 3.5. *Let P be a CCS process and Φ an $\text{LT}\mu$ formula. Then, $P \models \Phi$ if and only if $\Phi \preceq P$.*

Due to Prop. 3.4 and the definition of \models , it is sufficient to prove that $\llbracket \Phi \rrbracket = \mathcal{L}_{\text{max}}(\Phi) \cup \mathcal{L}_\omega(\Phi)$. This can be done along the structure of $\text{LT}\mu$ formulas, but requires the appropriate extension of the definition of languages to open terms.

3.4. Compositionality. One virtue of process algebras is that they allow for reasoning compositionally about processes. Our logical process calculus LPC is no exception. Indeed our must-preorder is compositional for all operators, except for the choice operators $+$ and \vee . This compositionality defect manifests itself in many behavioral preorders, including DeNicola and Hennessy’s must-preorder. The largest precongruence \sqsubseteq contained in \preceq can be obtained in the standard fashion [11].

DEFINITION 3.6 (Must-precongruence). *For $P, Q \in \mathcal{P}$ we write $P \sqsubseteq Q$ if (i) $P \preceq Q$ and (ii) $Q \xrightarrow{\tau} \text{ implies } P \xrightarrow{\tau}$.*

We can now establish the desired compositionality result.

THEOREM 3.7. *The preorder \sqsubseteq is a precongruence, i.e., for all processes P, Q such that $P \sqsubseteq Q$, the following properties hold:*

- $\alpha.P \sqsubseteq \alpha.Q$ for all $\alpha \in \mathcal{A}$
- $P + R \sqsubseteq Q + R$ for all $R \in \mathcal{P}$
- $P \vee R \sqsubseteq Q \vee R$ for all $R \in \mathcal{P}$
- $P|R \sqsubseteq Q|R$ for all $R \in \mathcal{P}$
- $P \wedge R \sqsubseteq Q \wedge R$ for all $R \in \mathcal{P}$
- $P \setminus L \sqsubseteq Q \setminus L$ for all restriction sets L
- $P[f] \sqsubseteq Q[f]$ for all relabelings f
- $\mu_k x.P \sqsubseteq \mu_k x.Q$ for all $x \in \mathcal{V}$ and $k \in \mathbb{N}$
- $\mu x.P \sqsubseteq \mu x.Q$ for all $x \in \mathcal{V}$
- $\nu x.P \sqsubseteq \nu x.Q$ for all $x \in \mathcal{V}$

Moreover, \sqsubseteq is the largest precongruence contained in \sqsubset .

The compositionality property can be checked straightforwardly for most operators by referring to Thm. 3.2. For asynchronous parallel composition, the compositionality of \sqsubseteq follows immediately from the fact that $P|Q \text{ must } T$ if and only if $P \text{ must } Q|T$, for all $P, Q \in \mathcal{P}$ and $T \in \mathcal{T}$; this is essentially the associativity property of $|$. The proof of the ‘largest’ statement of Thm. 3.7 is standard [11].

4. Discussion and Related Work. This section compares LPC to related work and discusses in some detail the fundamental differences of the setting presented in this paper to our previous approach [9].

Most early related work couples operational and declarative approaches to system specification loosely and does not allow for mixed specifications. This includes the large amount of work on relating behavioral equivalences or preorders to temporal logics in one of the following ways: (i) establishing that one system refines another if and only if both satisfy the same temporal formulas [12, 17, 25, 31]; (ii) translating finite-state labeled transition systems into temporal formulas [30]; or (iii) encoding subclasses of temporal formulas as behavioral relations via the idea of implicit specifications [23]. Other work, in the field of compositional model checking [8, 14, 20] is aimed at supporting a modular approach for reasoning about temporal-logic specifications. Several researchers have also considered the inclusion of different fixed-point operators in behavioral theories of processes in order to model fairness and unbounded but finite delay [15, 18]. One may also find a process algebra with an element similar to our process ff in [2].

Diverting from these approaches, advanced frameworks for genuine heterogeneous specifications have been developed as well, which can be distinguished whether they are logic/algebraic or automata-theoretic.

4.1. Logic/algebraic approaches. This category includes the seminal work of Abadi and Lamport, who have developed ideas for heterogeneous specifications for shared-memory systems [1]. Their technical setting is the logical framework of TLA [22], in which processes and temporal formulas are indistinguishable and logical implication serves as the refinement relation. The difference to our setting is that TLA refinement is insensitive to deadlock and divergence. While this might not be a problem for shared-memory systems, it is not suitable for reasoning about distributed systems, at which our calculus LPC aims. Graf and Sifakis follow a similar line in [13]. There, a logic is developed that includes constructs for actions and nondeterministic choice, and a logical encoding of operational behavior is given. One establishes that a system satisfies a property by showing that the logical formula associated with the system implies the property.

In a different line of research, Valmari et al. have studied several congruences preserving “next-time-less” linear-time temporal logic [27], which may also handle deadlock and livelock [19, 28, 33]. A good overview by Puhakka and Valmari on the matters of liveness and fairness in process algebra can be found in [29]. This paper also observes that, during system refinement, fairness constraints are often only relevant for intermediate systems and are automatically implied when considering the larger system context. It then suggests a way to avoid constructing the usually infinite intermediate systems. Our work complements theirs in that LPC allows for embedding arbitrary LTL formulas in operational specifications, instead of a specific

class of fairness constraints. However, **LPC** does not avoid reasoning about infinite intermediate systems, since we believe that such reasoning poses no problem when employing clever data structures for implementing our must-preorder in verification tools. Finally, note that DeNicola and Hennessy’s testing theory [11] has also been enriched with notions of fairness [6, 26] to constrain infinite computations in transition systems.

4.2. Automata-theoretic approaches. Regarding automata-theoretic techniques, the work of Kurshan [21], who presented a theory of ω -word automata that includes notions of synchronous and asynchronous composition, is of direct relevance to this paper. However, Kurshan’s underlying semantic model maps processes to their infinite traces, and the associated notion of refinement is (reverse) trace inclusion. In theories of concurrency, such as in ours in which deadlock is possible, maximal trace inclusion is not compositional [24].

The most closely related approach to the one presented here was introduced by the authors in [9]. Büchi automata were employed to uniformly encode mixed operational and declarative behavior, exploiting the well-known relation between Büchi automata and LTL [34]. We equipped this semantic framework with a notion of Büchi must-testing that extends DeNicola and Hennessy’s must-testing preorder from labeled transition systems to Büchi automata. The intuition was only to consider those infinite traces as infinite computations that go through Büchi states infinitely often, and only to accept those infinite computations for which the considered Büchi test declares success infinitely often. The relation of our Büchi must-preorder to the LTL satisfaction relation, with the central result intended to be analogous to Thm. 3.5, was then established in a pure automata-theoretic fashion by suitably adapting the construction of [34]. However, our previous approach had several shortcomings which made it unsuitable as a semantic basis for a logical process calculus; these are discussed next.

Most importantly, our paper [9] contained a subtle technical mistake in the analogue of Lemma 3.3, which propagated through the paper’s results. In a nutshell, the setup of Büchi testing did not allow us, as was intended, to ignore non-Büchi divergent traces, i.e., those infinite internal computations that go through Büchi states only finitely often. While most of the results of [9] could be repaired by explicitly observing non-Büchi divergence, the framework did no longer reflect the underlying intuition, and it made compositionality difficult to achieve for some operators, including parallel composition. Moreover, our identification of ff , or other inconsistent specifications, with non-Büchi divergence lead to the invalidity of the desired law $P \vee \text{ff} \approx P$. The present paper repairs this defect by associating ff with a process that cannot engage in any observable transition, nor in any divergence. In order to then distinguish ff from, say, $\mathbf{0}$ we introduced the unimplementability predicate. Similar difficulties arose when interpreting tt as Büchi-divergent process, which is why this paper distinguishes between tt and Ω , making tt the smallest *convergent* process with respect to our must-preorder, while Ω is still the smallest process overall.

Indeed, the collection of these insights also allowed us to do away with Büchi automata as our semantic framework for heterogeneous system design altogether. Accordingly, **LPC** encodes the least and greatest fixed-points occurring in temporal logics via labeled transition systems, where the process-algebraic semantic rules for least fixed-points reflect the intuition that the recursion under consideration can only be unwound finitely often, while a recursion associated with a greatest fixed-point may be unwound infinitely often. Hence, in **LPC** all infinite traces are ‘good’, which means that the expressive power of Büchi automata to distinguish ‘good’ and ‘bad’ infinite traces is no longer needed. The result is a process calculus, **LPC**, in which classical process algebras and linear-time temporal logics can be uniformly integrated, as was envisioned in [9]. The integration is mathematically elegant, as testified by our compositionality and conservative extension results that were established in a pure syntax-driven manner.

5. Example: Heterogeneous System Design. This section illustrates, by means of an example, the kind of refinement-based system design supported by LPC. The example advocates a heterogeneous style of system specification, combining process-algebraic and temporal-logic specifications, and thereby testifies to the utility of our calculus. It will be convenient to express temporal constraints by means of formulas in *Linear-time Temporal Logic* (LTL) [27] — a temporal logic that engineers often prefer over the linear-time μ -calculus [5]. We thus briefly show how LTL formulas can be encoded in $\text{LT}\mu$ or, more precisely, in our new calculus LPC.

5.1. Encoding of LTL in LPC. Since we would like to describe action-based distributed systems and their deadlock behavior, the variant of LTL studied here includes the atomic propositions a , for $a \in \mathcal{A}$, and $\mathbf{0}$. Note that, in the context of temporal logics, \mathcal{A} is always taken to be a finite set.

$$\Phi ::= \mathbf{0} \mid a \mid \text{tt} \mid \text{ff} \mid \Phi \vee \Phi \mid \Phi \wedge \Phi \mid X\Phi \mid \hat{X}\Phi \mid \Phi \text{U} \Phi \mid \Phi \text{V} \Phi$$

The temporal operators X , U , and V are intuitively interpreted as *next*, *until*, and *release* operators, respectively. Operator \hat{X} is the dual operator of X , which is a *next* operator that tolerates deadlocks; note that X is not self-dual in the presence of finite traces. An LTL formula Φ corresponds to the LPC process $\llbracket \Phi \rrbracket$, where the translation function $\llbracket \cdot \rrbracket$ is defined inductively along the structure of Φ as follows and where x is some randomly chosen variable in \mathcal{V} .

$$\begin{aligned} \llbracket \text{tt} \rrbracket &:= \text{tt} & \llbracket \mathbf{0} \rrbracket &:= \mathbf{0} & \llbracket \Phi_1 \vee \Phi_2 \rrbracket &:= \llbracket \Phi_1 \rrbracket \vee \llbracket \Phi_2 \rrbracket & \llbracket X\Phi \rrbracket &:= \bigvee_{a \in \mathcal{A}} a. \llbracket \Phi \rrbracket \\ \llbracket \text{ff} \rrbracket &:= \text{ff} & \llbracket a \rrbracket &:= a.\text{tt} & \llbracket \Phi_1 \wedge \Phi_2 \rrbracket &:= \llbracket \Phi_1 \rrbracket \wedge \llbracket \Phi_2 \rrbracket & \llbracket \hat{X}\Phi \rrbracket &:= \mathbf{0} \vee \bigvee_{a \in \mathcal{A}} a. \llbracket \Phi \rrbracket \\ \llbracket \Phi_1 \text{U} \Phi_2 \rrbracket &:= \mu x. \llbracket \Phi_2 \rrbracket \vee (\llbracket \Phi_1 \rrbracket \wedge \bigvee_{a \in \mathcal{A}} a.x) \\ \llbracket \Phi_1 \text{V} \Phi_2 \rrbracket &:= \nu x. \llbracket \Phi_2 \rrbracket \wedge (\llbracket \Phi_1 \rrbracket \vee \mathbf{0} \vee \bigvee_{a \in \mathcal{A}} a.x) \end{aligned}$$

For convenience, we abbreviate formula $\text{ffV}\Phi$ by $G\Phi$ (“*generally* Φ ”) and $\text{ttU}\Phi$ by $F\Phi$ (“*eventually* Φ ”), as usual. Moreover, we let $a \implies \Phi$ stand for the process $a.\Phi \vee \mathbf{0} \vee \bigvee_{a \neq b} b.\text{tt}$ that is valid if and only if, for all traces of the form aw , trace w satisfies Φ .

5.2. Example. Suppose an engineer is expected to design a reliable bidirectional network link in a component-based fashion. One might think of this link as a composition of two reliable unidirectional links that are closely tied together. In particular, the failure of one unidirectional link should imply the failure of the other, which is a typical physical constraint of bidirectional links. The engineer might begin with a simple specification of an unreliable unidirectional link,

$$\text{ULSpec} := \nu x. \overline{\text{up}}.(x + \overline{\text{fail}}.\nu y. \overline{\text{down}}.(y \vee x)),$$

which signals whether the link is up or down, or whether it just failed. In case of failure, the link tries to repair itself and, if and once it is successfully repaired, it returns to its initial state. However, a successful repair is not guaranteed, whence the process ULSpec may infinitely engage in the $\overline{\text{down}}$ -loop over variable y .

To obtain a specification RLSpec of a reliable unidirectional link, ULSpec is simply refined by adding a constraint imposing a “repair guarantee,” $\text{RG} := G(\text{fail} \implies F\overline{\text{up}})$, i.e., every broken link is eventually repaired and up. We then define $\text{RLSpec} := \text{ULSpec} \wedge \text{RG}$, which essentially does away with the $\overline{\text{down}}$ -loop

in ULSpec . The desired bidirectional link might then be specified as follows:

$$\begin{aligned} \text{BLSpec} := & (\text{RLSpec}[\text{up1/up, down1/down, sync/fail}] \\ & | \text{RLSpec}[\text{up2/up, down2/down, } \overline{\text{sync}}/\text{fail}] \\ &) \setminus \{\text{sync}\}, \end{aligned}$$

where the synchronization on action fail , via the relabeling to action sync , ensures that the failure of one unidirectional link implies the failure of the other. Note that the constraints RG indirectly refer to action sync , which is restricted in BLSpec .

The engineer may now refine the heterogeneous LPC specification BLSpec into a pure CCS implementation. The idea is to fulfill the constraints RG by eliminating the $\overline{\text{down}}$ -loop in ULSpec , thus encoding that a repair can always be successfully carried out immediately. The implementation of RLSpec might accordingly be chosen as the CCS process $\text{RLImp} := \nu x. \overline{\text{up}}.(x + \overline{\text{fail}}.\overline{\text{down}}.x)$. We now establish that RLImp indeed refines RLSpec in the framework of our must-precongruence. First of all, it is easy to see by our characterization of \sqsubseteq (cf. Thm. 3.2) that $\text{ULSpec} \sqsubseteq \text{RLImp}$, due to the internal *nondeterministic* choice in ULSpec . Further, we obviously have $\text{RLImp} \models \text{RG}$. Hence, we may infer by Thm. 3.5 that $\text{RG} \sqsubseteq \text{RLImp}$. Because RLImp cannot engage in an initial τ -transition, we may in summary conclude $\text{ULSpec} \sqsubseteq \text{RLImp}$ and $\text{RG} \sqsubseteq \text{RLImp}$. By Prop. 2.2, which is also valid for \sqsubseteq , and by Thm. 3.7, we derive $\text{RLSpec} \equiv \text{ULSpec} \wedge \text{RG} \sqsubseteq \text{RLImp} \wedge \text{RLImp} \sqsubseteq \text{RLImp}$, as desired.

When replacing in BLSpec the components RLSpec by RLImp we obtain an implementation of our reliable bidirectional link, to which we refer as BLImp . Since \sqsubseteq is a precongruence and $\text{RLSpec} \sqsubseteq \text{RLImp}$, we obtain $\text{BLSpec} \sqsubseteq \text{BLImp}$, i.e., BLImp refines BLSpec , which coincides with our intuition.

Finally, it is worth mentioning that LPC actually may be seen as a temporal logic that allows for some restricted form of branching-time reasoning. For example, the LPC process $\text{sync} \implies (\overline{\text{down1}}.\text{tt} + \overline{\text{down2}}.\text{tt})$ encodes the property that the system state reached when executing action sync has both actions $\overline{\text{down1}}$ and $\overline{\text{down2}}$ enabled. Observe that, in contrast to $\overline{\text{down1}}.\text{tt} + \overline{\text{down2}}.\text{tt}$, the term $\overline{\text{down1}}.\text{tt} \wedge \overline{\text{down2}}.\text{tt}$ in LPC specifies the obvious contradiction that every initial transition is labeled by both actions $\overline{\text{down1}}$ and $\overline{\text{down2}}$ at the same time.

6. Conclusions and Future Work. We presented a novel *logical process calculus* LPC that integrates both classical process calculi, such as Milner's CCS, and temporal logics, such as the alternation-free linear-time μ -calculus $\text{LT}\mu$. The syntax of LPC enriched CCS by operators for synchronous parallel composition (conjunction) and nondeterministic choice (disjunction), as well as by minimal fixed-points operators (finite unwindings of recursion). The semantics of LPC was given in terms of labeled transition systems and an unimplementability predicate, both defined via structural operational rules. A refinement preorder on process terms was then introduced, which conservatively extends both DeNicola's and Hennessy's must-preorder and the $\text{LT}\mu$ satisfaction relation. Hence, $\text{LT}\mu$ model checking may as well be understood as refinement checking. Finally, our must-preorder was also shown to be compositional with respect to all operators in LPC.

The outcome of our studies is a heterogeneous specification language, which allows system designers to specify systems in a mixed operational and declarative style, together with a behavioral preorder that permits component-based refinement. We believe that our setting provides groundwork for formally investigating those software engineering languages that support heterogeneous specifications as a mixture of operational state machines and declarative constraints, such as the *Unified Modeling Language* [4].

Regarding future work, we intend to study axiomatizations of our must-preorder. We also plan to develop an algorithm for computing the must-preorder with the goal of implementing LPC in automated verification tools, such as the *Concurrency Workbench NC* [10].

REFERENCES

- [1] M. ABADI AND L. LAMPORT, *Composing specifications*, TOPLAS, 15 (1993), pp. 73–132.
- [2] J. BAETEN AND J. BERGSTRA, *Process algebra with a zero object*, in CONCUR' 90, Vol. 458 of LNCS, Springer-Verlag, 1990, pp. 83–98.
- [3] J. BERGSTRA, A. PONSE, AND S. SMOLKA, *Handbook of Process Algebra*, Elsevier Science, 2001.
- [4] G. BOOCH, J. RUMBAUGH, AND I. JACOBSON, *The Unified Modeling Language User Guide*, Addison Wesley Longman, 1998.
- [5] J. BRADFIELD AND C. STIRLING, *Modal logics and mu-calculi: An introduction*, in Handbook of Process Algebra, Elsevier Science, 2001, pp. 293–330.
- [6] E. BRINKSMA, A. RENSINK, AND W. VOGLER, *Fair testing*, in CONCUR '95, Vol. 962 of LNCS, Springer-Verlag, 1995, pp. 313–328.
- [7] S. D. BROOKES AND A. W. ROSCOE, *An improved failures model for communicating processes*, in Seminar on Concurrency, Vol. 197 of LNCS, Springer-Verlag, 1984, pp. 281–305.
- [8] E. CLARKE, D. LONG, AND K. MCMILLAN, *Compositional model checking*, in LICS '89, IEEE Comp. Society Press, 1989, pp. 353–362.
- [9] R. CLEAVELAND AND G. LÜTTGEN, *A semantic theory for heterogeneous system design*, in FSTTCS 2000, Vol. 1974 of LNCS, Springer-Verlag, 2000, pp. 312–324.
- [10] R. CLEAVELAND AND S. SIMS, *The NCSU Concurrency Workbench*, in CAV '96, Vol. 1102 of LNCS, 1996, pp. 394–397.
- [11] R. DENICOLA AND M. HENNESSY, *Testing equivalences for processes*, TCS, 34 (1983), pp. 83–133.
- [12] R. DENICOLA AND F. VAANDRAGER, *Three logics for branching bisimulation*, J. of the ACM, 42 (1995), pp. 458–487.
- [13] S. GRAF AND J. SIFAKIS, *A logic for the description of non-deterministic programs and their properties*, Information and Control, 68 (1986), pp. 254–270.
- [14] O. GRUMBERG AND D. LONG, *Model checking and modular verification*, TOPLAS, 16 (1994), pp. 843–871.
- [15] C. HARTONAS, *A fixpoint approach to finite delay and fairness*, TCS, 198 (1998), pp. 131–158.
- [16] M. HENNESSY, *Algebraic Theory of Processes*, MIT Press, 1988.
- [17] M. HENNESSY AND R. MILNER, *Algebraic laws for nondeterminism and concurrency*, J. of the ACM, 32 (1985), pp. 137–161.
- [18] T. HILDEBRANDT, *A fully abstract presheaf semantics of SCCS with finite delay*, in CTCS '99, Vol. 29 of ENTCS, Elsevier Science, 1999.
- [19] R. KAIVOLA AND A. VALMARI, *The weakest compositional semantic equivalence preserving nexttime-less linear temporal logic*, in CONCUR '92, Vol. 630 of LNCS, Springer-Verlag, 1992, pp. 207–221.
- [20] O. KUPFERMAN AND M. VARDI, *Modular model checking*, in Compositionality: The Significant Difference, Vol. 1536 of LNCS, Springer-Verlag, 1997.
- [21] R. KURSHAN, *Computer-Aided Verification of Coordinating Processes: The Automata-Theoretic Approach*, Princeton Univ. Press, 1994.

- [22] L. LAMPORT, *The temporal logic of actions*, TOPLAS, 16 (1994), pp. 872–923.
- [23] K. LARSEN, *The expressive power of implicit specifications*, TCS, 114 (1993), pp. 119–147.
- [24] M. MAIN, *Trace, failure and testing equivalences for communicating processes*, J. of Par. Comp., 16 (1987), pp. 383–400.
- [25] R. MILNER, *Communication and Concurrency*, Prentice Hall, 1989.
- [26] V. NATARAJAN AND R. CLEAVELAND, *Divergence and fair testing*, in ICALP '95, Vol. 944 of LNCS, Springer-Verlag, 1995, pp. 684–695.
- [27] A. PNUELI, *The temporal logic of programs*, in FOCS '77, IEEE Comp. Society Press, 1977, pp. 46–57.
- [28] A. PUHAKKA AND A. VALMARI, *Weakest-congruence results for livelock-preserving equivalences*, in CONCUR '99, Vol. 1664 of LNCS, Springer-Verlag, 1999, pp. 510–524.
- [29] ———, *Liveness and fairness in process-algebraic verification*, in CONCUR 2001, Vol. 2154 of LNCS, Springer-Verlag, 2001, pp. 202–217.
- [30] B. STEFFEN AND A. INGÓLFSDÓTTIR, *Characteristic formulae for CCS with divergence*, Inform. & Comp., 110 (1994), pp. 149–163.
- [31] C. STIRLING, *Modal logics for communicating systems*, TCS, 49 (1987), pp. 311–347.
- [32] ———, *Modal and Temporal Properties of Processes*, Springer-Verlag, 2001.
- [33] A. VALMARI AND M. TIERNARI, *Compositional failure-based semantics models for basic LOTOS*, FAC, 7 (1995), pp. 440–468.
- [34] M. VARDI AND P. WOLPER, *An automata-theoretic approach to automatic program verification*, in LICS '86, IEEE Comp. Society Press, 1986, pp. 332–344.

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.				
1. AGENCY USE ONLY(Leave blank)	2. REPORT DATE August 2002	3. REPORT TYPE AND DATES COVERED Contractor Report		
4. TITLE AND SUBTITLE A LOGICAL PROCESS CALCULUS		5. FUNDING NUMBERS C NAS1-97046 WU 505-90-52-01		
6. AUTHOR(S) Rance Cleaveland and Gerald Lüttgen				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) ICASE Mail Stop 132C NASA Langley Research Center Hampton, VA 23681-2199		8. PERFORMING ORGANIZATION REPORT NUMBER ICASE Report No. 2002-26		
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) National Aeronautics and Space Administration Langley Research Center Hampton, VA 23681-2199		10. SPONSORING/MONITORING AGENCY REPORT NUMBER NASA/CR-2002-211759 ICASE Report No. 2002-26		
11. SUPPLEMENTARY NOTES Langley Technical Monitor: Dennis M. Bushnell Final Report Submitted to the Ninth International Workshop on Expressiveness in Concurrency (EXPRESS 2002).				
12a. DISTRIBUTION/AVAILABILITY STATEMENT Unclassified-Unlimited Subject Category 60, 61 Distribution: Nonstandard Availability: NASA-CASI (301) 621-0390		12b. DISTRIBUTION CODE		
13. ABSTRACT (Maximum 200 words) This paper presents the Logical Process Calculus (LPC), a formalism that supports heterogeneous system specifications containing both operational and declarative subspecifications. Syntactically, LPC extends Milner's Calculus of Communicating Systems with operators from the alternation-free linear-time μ -calculus ($LT\mu$). Semantically, LPC is equipped with a behavioral preorder that generalizes Hennessy's and DeNicola's must-testing preorder as well as $LT\mu$'s satisfaction relation, while being compositional for all LPC operators. From a technical point of view, the new calculus is distinguished by the inclusion of (i) both minimal and maximal fixed-point operators and (ii) an unimplementability predicate on process terms, which tags inconsistent specifications. The utility of LPC is demonstrated by means of an example highlighting the benefits of heterogeneous system specification.				
14. SUBJECT TERMS heterogeneous specification, must-testing, process algebra, temporal logic, testing theory			15. NUMBER OF PAGES 19	
			16. PRICE CODE A03	
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT	